

# The Hitchhiker’s Guide to Efficient, End-to-End, and Tight DP Auditing

Meenatchi Sundaram Muthu Selva Annamalai<sup>1</sup>, Borja Balle<sup>2</sup>, Jamie Hayes<sup>2</sup>,  
Georgios Kaissis<sup>2</sup>, and Emiliano De Cristofaro<sup>3</sup>

<sup>1</sup>University College London, <sup>2</sup>Google DeepMind, <sup>3</sup>UC Riverside

## Abstract

This paper systematizes research on auditing Differential Privacy (DP) techniques, aiming to identify key insights into the current state of the art and open challenges. First, we introduce a comprehensive framework for reviewing work in the field and establish three cross-contextual desiderata that DP audits should target—namely, efficiency, end-to-end-ness, and tightness. Then, we systematize the modes of operation of state-of-the-art DP auditing techniques, including threat models, attacks, and evaluation functions. This allows us to highlight key details overlooked by prior work, analyze the limiting factors to achieving the three desiderata, and identify open research problems. Overall, our work provides a reusable and systematic methodology geared to assess progress in the field and identify friction points and future directions for our community to focus on.

## 1 Introduction

In today’s data-driven economy, a large number of potentially sensitive datasets are collected, shared, and released, including demographic [106], behavioral [105], and medical [85] data. Alas, this comes with inherent privacy risks that are difficult to mitigate in practice [81, 102]. In this context, Differential Privacy (DP) [39] provides a robust mathematical framework offering formal privacy guarantees. More precisely, DP bounds the impact that any single user can have on the output of a data release. In recent years, DP techniques have been deployed in both the public and private sectors. High-profile examples include the 2020 US Census release [2], Apple’s local DP system for learning user preferences [9], Google’s next-word prediction models [74], and Microsoft’s synthetic data release with the UN’s International Organization for Migration [77]. (For a comprehensive and up-to-date list of real-world DP deployments, we refer the reader to [32].)

**DP Auditing.** In practice, implementing differentially private algorithms correctly can be quite challenging. In the presence of incorrect implementations and/or bugs, the guarantees provided by DP can be severely degraded [6, 82] or completely compromised [56, 104], resulting in the unintended leakage of sensitive user data. This has prompted extensive research on verifying whether DP’s theoretical

guarantees are met in practice [8, 19, 34, 54, 66, 82, 100, 111], a process known as *DP auditing* [35]. As outlined in Figure 1, this involves conducting experiments on a mechanism’s implementation, analyzing its outputs, and comparing against the theoretical guarantees of the mechanism’s *specification*.

DP auditing is also a powerful and versatile technique to investigate the *tightness* of theoretical guarantees [7, 54, 71, 82, 83], develop theoretical intuitions for improved privacy analyses [20, 84], and refute conjectures [5]. Moreover, it can be used to bound the adversarial success with respect to privacy attacks like membership inference [54], attribute inference [72], or reconstruction [51, 71].

**Motivation.** To the best of our knowledge, research on DP auditing dates back to 2013 [35]; yet, extending and improving DP auditing techniques or proposing new ones remains an active area of research today. Early work [19, 34] required running the implementations millions of times, thereby making their computational overhead prohibitively high for auditing complex private machine learning pipelines. More recent efforts have focused on *efficiency*, aiming to reduce the number of implementation runs down to a single one [71, 100, 111]. However, for complex DP mechanisms, these techniques are still largely ineffective in identifying bugs and violations as the audits fail to produce estimates of the empirical privacy leakage that are “close” to the theoretical DP bounds—i.e., the audits are not “tight.”

While some results, e.g., [82], provide audits that are both efficient and tight, they require aggressive modifications to the implementations, including adversarial assumptions that might not always be realistic in practice. Thus, these techniques may fail to detect a large class of bugs or extend to other mechanisms; we refer to this characteristic as not being “end-to-end.” This prompts the need for a unified approach to building and evaluating DP auditing techniques that can simultaneously be efficient, end-to-end, and tight, arguably achieving the “gold standard” of DP auditing.

**Roadmap.** In this paper, we set out to systematize the large body of recent work on DP auditing along several key axes, aiming to identify fresh insights, open research problems, and promising directions to guide future work in improving auditing techniques. We define and formalize three desiderata for DP auditing – efficient, end-to-end, and tight – and provide a comprehensive framework for evaluating progress in the field. Finally, we set out to systematize the modes

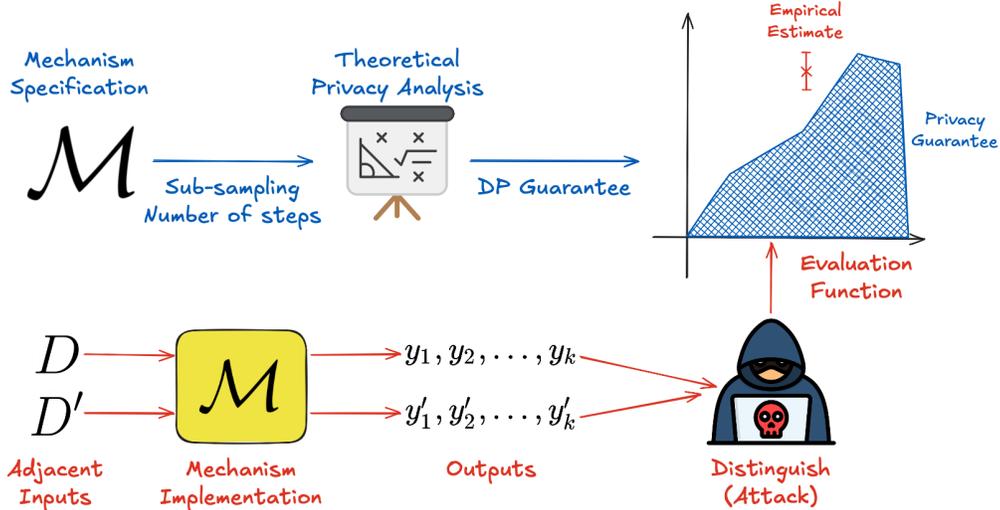


Figure 1: High-Level Overview of DP Auditing.

of operation of current DP auditing techniques, including threat models, attacks, and evaluation functions.

Our analysis sheds light on key details overlooked by prior work and the limiting factors in achieving the three desiderata. For instance, one of the main limitations to tighter audits is the lack of tight theoretical privacy analyses for sub-sampling schemes, such as shuffling. We also show that, while recent audits have focused on the restrictive “black-box” threat model (where adversaries can only insert a single input sample and observe only the algorithm’s final output), much less attention has been given to slightly stronger threat models where audits are equally not tight. Finally, we find that, while recent work achieves efficient and tight audits for simple mechanisms, this remains far from reach, even in strong threat models, for more realistic and complex mechanisms like DP-SGD [1].

**Contributions.** In summary, our work makes three key contributions:

- We introduce a comprehensive framework to systematize state-of-the-art research on DP auditing with respect to its key aspects, namely the algorithmic foundations and the modes of operation involved.
- We analyze the progress made by work in this field vis-à-vis three main desired properties, i.e., efficiency, end-to-end-ness, and tightness, underscoring the importance of achieving all three simultaneously both for simple mechanisms and complex private machine learning pipelines.
- We discuss several insights and directions that future work could focus on to improve auditing techniques.

## 2 Background

In this section, we review background concepts related to Differential Privacy. Readers familiar with these topics can skip it without losing continuity.

### 2.1 Differential Privacy

Differential Privacy (DP) provides a rigorous framework to control the privacy risk to individuals in a dataset. It ensures that the output of a computation, e.g., an answer to a statistical query or inference from a machine learning model, is not significantly affected by the presence/absence of any one individual’s data—commensurate to a parameter  $\epsilon$ . Over the years, a few variants of the DP definition have been proposed, as reviewed next.

**Definition 2.1 (Pure DP [37]).** A randomized mechanism  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  satisfies  $\epsilon$ -DP if, for any two adjacent datasets  $D, D' \in \mathcal{D}$  and  $S \subseteq \mathcal{R}$ , it holds:

$$\Pr[\mathcal{M}(D) \in S] \leq e^\epsilon \Pr[\mathcal{M}(D') \in S]$$

Pure DP is defined by a single  $\epsilon$  parameter describing the upper bound on the probability an adversary can distinguish between outputs on two adjacent inputs (cf. Section 2.2).

**Approximate DP.**  $(\epsilon, \delta)$ -DP is relaxation of Pure DP that introduces a “failure probability”  $\delta$ , i.e., the probability that the  $\epsilon$ -DP guarantee does not hold.

**Definition 2.2 (Approximate DP [38]).** A randomized mechanism  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  satisfies  $(\epsilon, \delta)$ -DP if, for any two adjacent datasets  $D, D' \in \mathcal{D}$  and  $S \subseteq \mathcal{R}$ , it holds:

$$\Pr[\mathcal{M}(D) \in S] \leq e^\epsilon \Pr[\mathcal{M}(D') \in S] + \delta$$

**Rényi DP (RDP).**  $(\alpha, \gamma)$ -RDP captures a collection of  $(\epsilon, \delta)$ -DP guarantees in one definition to more tightly analyze privacy guarantees of DP mechanisms.

**Definition 2.3 (Rényi divergence [94]).** For any two probability distributions,  $P$  and  $Q$ , the Rényi divergence of order  $\alpha > 1$  is:

$$D_\alpha(P, Q) \triangleq \frac{1}{\alpha - 1} \log \mathbb{E}_{x \sim Q} \left( \frac{P(x)}{Q(x)} \right)^\alpha$$

$$\text{and } D_1(P, Q) = \lim_{\alpha \rightarrow 1} D_\alpha(P, Q) = \mathbb{E}_{x \sim P} \log \left( \frac{P(x)}{Q(x)} \right).$$

**Definition 2.4** (Rényi DP [79]). *A randomized mechanism  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  satisfies  $(\alpha, \gamma)$ -RDP if, for any two adjacent datasets  $D, D' \in \mathcal{D}$ , it holds:*

$$D_\alpha(\mathcal{M}(D), \mathcal{M}(D')) \leq \gamma$$

where  $D_\alpha$  is the Rényi divergence of order  $\alpha \geq 1$ .

**Functional DP ( $f$ -DP).**  $f$ -DP uses the hypothesis testing interpretation of DP [58] (reviewed in Section 2.4) to bound the type II error achievable by an adversary aiming to distinguish between the outputs of a DP mechanism on adjacent inputs at a significance level of  $\alpha$ . A specific instantiation of  $f$ -DP is Gaussian Differential Privacy (aka  $\mu$ -GDP), for  $f = T(\mathcal{N}(0, 1), \mathcal{N}(\mu, 1))$ .

**Definition 2.5** (Trade-off function [36]). *For any two probability distributions on the same space,  $P$  and  $Q$ , the trade-off function  $T(P, Q) : [0, 1] \rightarrow [0, 1]$  is defined as:*

$$T(P, Q)(\alpha) \triangleq \inf_{\phi} \{\beta_{\phi} : \alpha_{\phi} \leq \alpha\}$$

with the infimum taken over all rejection rules  $\phi$  for which  $\alpha_{\phi}$  and  $\beta_{\phi}$  are the type I and type II errors, respectively.

**Definition 2.6** (Functional DP [36]). *A randomized mechanism  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  satisfies  $f$ -DP if, for any two adjacent datasets  $D, D' \in \mathcal{D}$ , and  $\alpha \in [0, 1]$  it holds:*

$$T(\mathcal{M}(D), \mathcal{M}(D'))(\alpha) \geq f(\alpha).$$

## 2.2 Adjacency

In the above DP definitions, the notion of adjacent (aka neighboring) datasets is deliberately left generic as it varies depending on the setting. We now review three common notions of adjacency:

- **Add/Remove:** corresponds to inserting or deleting a record from the dataset. It is also referred to as “unbounded DP” since the dataset size is not the same, i.e.,  $|D| = |D'| \pm 1$ .
- **Edit:** replacing one record with another. It is also referred to as “bounded DP”, i.e.,  $|D| = |D'|$ . Guarantees under edit adjacency are typically twice as strong as add/remove as one edit corresponds to one remove and one add.
- **Zero-out [57]:** introduced to bridge guarantees provided under the add/remove and edit adjacencies and simplify theoretical privacy analyses. Two datasets are zero-out adjacent if exactly one record in one dataset is replaced with a special *zero-out record*,  $\perp$  in the other.

## 2.3 Sub-Sampling

A common technique used to amplify a mechanism’s DP guarantees and/or make DP compatible with modern machine-learning techniques is known as *sub-sampling*. More precisely, a mechanism is run iteratively on small, randomly chosen subsets of the input dataset, resulting in

substantially improved privacy guarantees compared to running the mechanism on the full dataset. This is referred to as “privacy amplification by sub-sampling” [22] and depends on the specific sub-sampling scheme used. We now review four common ones:

- **Poisson:** at each iteration, each record is chosen with some sampling probability  $q$ .
- **Sampling w/ (or w/o) replacement:** a batch of  $B$  records is chosen, at each iteration, by sampling uniformly at random w/ (or w/o) replacement. For brevity, we denote these schemes as Sampling WR and WOR, respectively.
- **Shuffling:** the dataset is randomly permuted (shuffled) and then partitioned into  $b$  batches of  $B$  records. The batches are iterated over in sequence until the final one is processed, and the dataset is randomly permuted again.
- **Balls-in-bins [25, 27]:** each record is uniformly assigned to one of  $b$  batches, then, the batches are processed in sequence for each iteration. After the final batch is processed, the first one is processed again in a round-robin fashion.

## 2.4 Hypothesis Testing Interpretation of DP

The notion of hypothesis testing interpretation of DP, introduced in [58] for  $(\epsilon, \delta)$ -DP, allows to evaluate a mechanism’s privacy guarantees based on the difficulty an adversary faces in distinguishing between two adjacent datasets. Given a random output from a mechanism  $y \leftarrow \mathcal{M}(D^*)$  (where  $D^* \in \{D, D'\}$  for adjacent  $D, D'$ ), DP can be interpreted as inducing the following a hypothesis test:

$$H_0 : y \sim \mathcal{M}(D) \quad H_1 : y \sim \mathcal{M}(D')$$

**Privacy Region.** For a choice of rejection region  $S$ , i.e.,  $H_0$  is rejected if  $y \in S$ , the type I and type II errors are defined as  $P_{\text{FA}}(D, D', \mathcal{M}, S) = \Pr[\mathcal{M}(D) \in S]$  and  $P_{\text{MD}}(D, D', \mathcal{M}, S) = \Pr[\mathcal{M}(D') \notin S]$ , respectively. Thus,  $(\epsilon, \delta)$ -DP bounds the type I and type II errors as per the following theorem:

**Theorem 2.1** ([58]). *A randomized mechanism  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  satisfies  $(\epsilon, \delta)$ -DP if and only if the following conditions are satisfied for all adjacent datasets  $D, D'$  and all rejection regions  $S \subseteq \mathcal{R}$ :*

$$P_{\text{FA}}(D, D', \mathcal{M}, S) + e^{\epsilon} P_{\text{MD}}(D, D', \mathcal{M}, S) \geq 1 - \delta \wedge e^{\epsilon} P_{\text{FA}}(D, D', \mathcal{M}, S) + P_{\text{MD}}(D, D', \mathcal{M}, S) \geq 1 - \delta$$

This interpretation can also be operationalized, i.e., an analyst can pick a pair of adjacent datasets  $D, D'$  and plot the type I and type II errors obtained for different rejection regions  $S$ . This will form a “privacy region” defined by  $(\epsilon, \delta)$ -DP, which is often used for DP auditing.

## 3 DP Auditing

We now discuss the notion of DP auditing – a set of empirical techniques involving an adversary instantiated against

a DP mechanism to distinguish between the mechanism’s outputs on adjacent datasets. We discuss its two main motivating use cases and the desired properties we believe DP auditing algorithms should satisfy.

### 3.1 Goals of DP Auditing

Arguably, the two main goals of DP auditing include: 1) *verifying* that the theoretical guarantees provided by DP are met in practice and 2) *estimating* the actual privacy leakage observed from the mechanism empirically.

#### 3.1.1 Verifying Privacy Guarantees

While the increased adoption of DP in real-world applications [2, 9, 63, 74, 77, 87] is a promising development for privacy-enhancing technologies, it can often be difficult to implement DP correctly. Bugs leading to unintended privacy leakage, in some cases completely breaking DP guarantees, have been found in many implementations [6, 8, 47, 56, 66, 69, 117]. For instance, Lyu et al. [69] and Zhang et al. [117] focus on the differentially private Sparse Vector Technique (SVT) [40] and show, through manual analysis, that most of its implementations [23, 41, 64, 69, 95, 101] are not differentially private at all.

Bugs in the implementation of machine-learning algorithms like Differentially Private Stochastic Gradient Descent (DP-SGD) [1, 6, 56] and DP synthetic data [8, 47] have also been found through both manual and automated techniques. However, manual analysis can be error-prone and may miss subtle implementation details, thus prompting the need to scale up automated approaches [8].

**Limitations of Static Analysis.** Early work aiming to automatically verify DP guarantees in algorithms’ implementations has used static code analysis tools, e.g., type checkers and specialized programming languages. For example, Reed and Pierce [93] present a type system that enables programmers to write privacy-safe programs by design. Overall, a large body of work [3, 14–16, 46, 108, 109, 116] has focused on making type systems more expressive and efficient.

However, this approach faces some fundamental limitations. First, type systems implicitly trust that some “basic” DP algorithms (e.g., Gaussian mechanism, Laplace mechanism) are implemented correctly by the underlying language. As pointed out in [66], this may not necessarily be the case if, e.g., the noise addition mechanisms suffer from floating-point vulnerabilities [78]. Second, these type systems require mechanisms to be implemented in a very specific language, thus, they are substantially limited to simple algorithms like SVT or Report Noisy Max [41]. Specifically, they struggle to capture the complexities involved in algorithms like RAPPOR [43] or DP-SGD [1].

**Enter DP Auditing.** Rather than statically checking program specifications, DP auditing allows an auditor to run experiments and *statistically* refute the theoretical privacy guarantees provided by DP. In other words, if the auditor

fails, programmers and clients can be assured that the implementation satisfies the DP guarantees at some (possibly pre-defined) level of confidence—e.g., for more critical applications, the auditor can perform the test at 99.999% confidence [104] or 95% for less critical applications [54].

Overall, DP audits can offer more flexibility than static methods, accommodating more complex algorithms such as private machine learning pipelines [54]. They are also easier to implement as they do not involve specific programming languages or typing rules.

#### 3.1.2 Empirical Privacy Estimation

DP auditing can also be useful to provide precise estimates of the privacy leakage in different settings, e.g., vis-à-vis different real-world adversaries and/or attacks.

**From Worst-Case to Real-World Adversaries.** DP provides formal privacy guarantees against powerful theoretical adversaries that are conceivably not always realistic in practice. In some cases, these strong adversarial assumptions may be needed to provide robust protections that hold in worst-case settings. In other cases, this may instead be a direct consequence of limited known ways of proving privacy guarantees [83]. For instance, the privacy analysis of DP-SGD [1] assumes an adversary that has access not only to the final trained model but also to all intermediate models at each training step. This assumption is not entirely realistic across the board as intermediate parameters are not usually released, but needed since theoretical privacy analysis capabilities in this threat model are limited [26, 44, 112]. Therefore, DP auditing can be used to empirically estimate privacy leakage in real-world settings [5, 7, 20, 54, 84].

In this context, an empirical estimate  $\epsilon_{\text{emp}}$  can be derived experimentally and then compared against the theoretical bound  $\epsilon$ . This allows to investigate the *tightness* of theoretical guarantees [7, 54, 71, 82, 83], develop theoretical intuitions for improved privacy analyses [20, 84], and refute widely believed conjectures [5].

**Testing Resistance to Different Privacy Attacks.** Another use-case of DP auditing stems from its modularity. DP auditing can be instantiated, e.g., on different adjacency definitions or DP variants or using different attacks. While the DP definition directly maps to the notion of membership inference (MI), which we review in Section 5.2, DP guarantees can also be extended to and empirically verified against other attacks such as attribute inference (AI) and reconstruction attacks (Recon) [51, 96, 113]. Privacy analyses under more general privacy attacks may not always be optimal, thus, DP auditing can be used to investigate the tightness of the protections provided by DP against them [51, 71].

### 3.2 DP Auditing Algorithms

We now present a formal definition of a DP auditing algorithm, which we use in our systematization to identify relevant papers in this area.

---

**Algorithm 1** Prototypical Auditing Algorithm

---

**Require:** Mechanism,  $\mathcal{M}$ . DP Guarantee,  $\kappa \in \mathcal{K}$ . Data domain,  $\mathcal{X}$ . Number of runs,  $n_R$ . Number of canaries,  $n_C$ . Confidence level,  $\alpha$ . Privacy attack,  $\mathcal{P}$ . Evaluation function,  $\zeta$ .

▷ **Adversary**

- 1: **Pick dataset:**  $D^- \in \mathcal{X}^*$
- 2: **Pick canary distributions:**  $\mathcal{C}$
- 3: **Define auxiliary information:**  $\text{aux}$

▷ **Challenger**

- 4:  $Y = []$
- 5:  $D_C = []$
- 6: **for**  $i \in [n_R]$  **do**
- 7:      $D_C[i] \leftarrow \{x \sim \mathcal{C}[i] \mid i \in [n_C]\}$
- 8:      $Y[i] \leftarrow \mathcal{M}(D^- \cup D_C[i]; \text{aux})$
- 9: **end for**

▷ **Adversary**

- 10:  $O = []$
- 11: **for**  $i \in [n_R]$  **do**
- 12:      $O[i] \leftarrow \mathcal{P}(Y[i]; \mathcal{M}, D^-, \mathcal{C})$
- 13: **end for**

- 14: **return**  $z \leftarrow \zeta(O; \mathcal{M}, D_C, \mathcal{C}, \alpha, \kappa)$
- 

To ease our definition, we first define an ordering over the space of possible DP guarantees  $\mathcal{K}$ :

**Definition 3.1** (Ordering of DP Guarantees). *For  $\kappa, \kappa' \in \mathcal{K}$ ,  $\kappa \geq \kappa'$  if  $\mathcal{M}$  satisfies  $\kappa'$ -DP  $\rightarrow$   $\mathcal{M}$  satisfies  $\kappa$ -DP.*

In Pure DP,  $\mathcal{K} = \mathbb{R}^+ \cup \{0\}$ ,  $\kappa = \epsilon$ , and the ordering follows the natural ordering of real numbers.

**Definition 3.2** (DP Auditing Algorithm). *A DP auditing algorithm  $\mathcal{A}_{\text{DP}}$  takes in input a  $\kappa$ -DP mechanism  $\mathcal{M}$ , the claimed DP guarantees  $\tilde{\kappa} \in \mathcal{K}$ , and some additional hyperparameters  $\Theta$  and outputs  $z \leftarrow \mathcal{A}_{\text{DP}}(\mathcal{M}, \tilde{\kappa}; \Theta)$ .*

*If  $\mathcal{A}_{\text{DP}}$  is used for verification,  $z \in \{0, 1\}$  and  $\mathcal{A}_{\text{DP}}(\mathcal{M}, \tilde{\kappa}; \Theta) = 1$  if  $\tilde{\kappa} \geq \kappa$ . If used for empirical estimation,  $z \in \mathcal{K}$  and  $\kappa \geq \mathcal{A}_{\text{DP}}(\mathcal{M}, \tilde{\kappa}; \Theta)$ .*

This definition also encodes a notion of “correctness” for DP auditing algorithms, which informally states that the output of the algorithm agrees with the actual guarantees satisfied by the mechanism (i.e., if  $\mathcal{M}$  satisfies  $\kappa$ -DP, the algorithm confirms this to be true as well). However, this definition admits vacuous algorithms (e.g., an algorithm that always outputs ‘1’); thus, in Section 3.3, we define additional desirable properties of DP auditing algorithms.

**Auditing Game.** A DP mechanism is typically audited vis-à-vis a *privacy game* played between an *Adversary* and a *Challenger*. Although this game can slightly vary depending on the auditing goal and threat model, we present the prototypical DP audit algorithm in Algorithm 1.

First, the Adversary chooses an initial dataset in the data domain  $D^-$  and distributions from which special data points used to audit the mechanism, i.e., *canaries*, are sampled from. The Adversary can also provide additional “auxil-

ary” information to the mechanism if necessary. Next, the Challenger runs the mechanism  $n_R$  times, each time sampling  $n_C$  fresh canaries from the canary distributions to form the input dataset. The Challenger then sends the outputs back to the Adversary. The Adversary post-processes the outputs of the mechanism, typically using a privacy attack  $\mathcal{P}$ . This is because the outputs of the mechanism may be in an arbitrary domain  $y \in \mathcal{Y}$ , which can be difficult to design a decision function around (e.g., to decide if the mechanism satisfies the claimed guarantees). Thus, the post-processing yields simple observations that can be easily analyzed.

Finally, the audit’s outcome can be computed from the observations using an evaluation function  $\zeta$ , which outputs either a bit  $z \in \{0, 1\}$  for the verification of a theoretical DP guarantee or  $z \in \mathcal{K}$  for empirical privacy estimation. This is typically done by relying on the hypothesis testing interpretation of DP discussed above. In practice, the outcome of the audit is typically computed with a desired level of confidence  $\alpha$ , especially when refuting/confirming the claimed DP guarantees of the mechanism.

### 3.3 Desiderata: The Gold Standard of DP Auditing

Next, we define three key properties we believe audits should satisfy – *efficiency*, *end-to-end-ness*, *tightness*. We believe these properties are crucial for the widespread and rigorous adoption of DP auditing, especially in production systems. As discussed later, achieving all three simultaneously mostly remains an open research problem.

**Efficiency.** For DP audits to be deployable in practice, they have to be computationally efficient. Early work [18, 34] required implementations to be executed millions of times – while this is achievable for simple mechanisms (e.g., Laplace and Gaussian mechanisms), it can be prohibitively expensive for complex mechanisms like DP-SGD [1]. Therefore, recent work [4, 71, 100, 111] has focused on developing auditing techniques that only require implementations to be executed a handful of times, possibly even just once.

**Definition 3.3** (Efficient Audit). *A DP auditing algorithm  $\mathcal{A}_{\text{DP}}(\mathcal{M}, \kappa; \Theta)$  is efficient if it requires the mechanism  $\mathcal{M}$  to be run only a few times (typically once or twice).*

**End-to-End.** DP audits should also apply to *end-to-end* privacy guarantees of the implementations; ideally, these should be run by simply providing inputs to the mechanism and observing its final output. This is often referred to as “black-box” auditing [65] and reflects realistic threat models encountered in practical settings as opposed to worst-case threat models assumed by the privacy analysis of certain mechanisms. For instance, the privacy analysis of the DP-SGD [1] algorithm assumes that all intermediate models are released, even though in practice only the final model may be accessible to the adversary. Since “black box” may have different meanings in different areas (e.g., query-access to machine learning models [99] vs. access to fixed generated

synthetic data records [52], etc.), we use the notion of “end-to-end” auditing to avoid confusion.

End-to-end audits are also broadly applicable to many mechanisms as they do not require access to *how* they are implemented internally. For example, although different DP mechanisms can be used to train private ML models (e.g., DP-SGD [1], PATE [89], DP-ZO [103]), relying on one single auditing technique that takes in input the training dataset and observe the trained model’s outputs would make it significantly easier to integrate auditing into different training pipelines. Once again, this would also more closely match the capabilities of realistic adversaries.

**Definition 3.4** (End-to-End Audit). *A DP auditing algorithm  $\mathcal{A}_{\text{DP}}(\mathcal{M}, \kappa; \Theta)$  is end-to-end if it does not modify the mechanism  $\mathcal{M}$  and performs the audit simply by observing the mechanisms output(s) on specific input(s).*

**Tightness.** Informally, a *tight* audit means that the empirical audit is consistent with the theoretical privacy analysis. More precisely, the empirical estimates obtained through auditing should not be far from the theoretical guarantees (e.g.,  $\epsilon_{\text{emp}} \approx \epsilon$  and  $\delta_{\text{emp}} \approx \delta$ ).

Moreover, when verifying DP guarantees, tight DP audits would ensure that implementations are successfully verified only if their actual privacy leakage is strictly lower than the theoretical DP bounds. As a consequence, tight audits are needed in this context to reinforce their accuracy and provide strong confidence in the process. By contrast, a *loose* audit that successfully verifies implementations where the actual privacy leakage is, say, twice the theoretical upper bounds would inherently not be very useful to identify bugs.

**Definition 3.5** (Tight Audit). *A DP auditing algorithm  $\mathcal{A}_{\text{DP}}(\mathcal{M}, \kappa; \Theta)$  is tight if it it never underestimates the true privacy parameters, i.e., assuming  $\mathcal{M}$  satisfies  $\kappa$ -DP,  $\mathcal{A}_{\text{DP}}(\mathcal{M}, \kappa; \Theta) = \kappa$  for estimation and  $\forall \kappa' < \kappa$   $\mathcal{A}_{\text{DP}}(\mathcal{M}, \kappa'; \Theta) = 0$  for verification.*

## 4 The DP Auditing Body of Work

In this section, we present the methodology we use to compile a comprehensive and representative snapshot of research on DP auditing.

**Querying Scopus.** In January 2025, we queried the Scopus research database [42] using several keywords (see Appendix A for the complete list) and obtained a total of 179 papers that mention DP auditing or similar terms in their title or abstract. We use Scopus as it indexes the top-four Security and Privacy venues (IEEE S&P, ACM CCS, USENIX, and NDSS), the major Machine Learning conferences (e.g., NeurIPS, ICLR, ICML), as well as prominent journals like the Transactions of Machine Learning Research (TMLR).

**Manual Review.** We then manually reviewed the abstracts of the 179 articles and filtered relevant ones that either explicitly estimate the DP parameters of an implementation

or verify whether claimed DP guarantees are met following Definition 3.2. We also followed a snowballing approach (aka citation chaining) by examining the references cited in these papers to identify recently accepted work on DP auditing not yet indexed by Scopus, as well as non-peer-reviewed papers that already informed published research at top conferences (e.g., articles published on arXiv). This yielded an additional 24 papers.

**Final List.** After manual review, we identified 45 relevant DP auditing papers, which form the basis for our systematization effort. To ease presentation, in Table 1, we report 21 key papers (out of the 45) that arguably represent the state of the art, while reporting earlier/seminal work in Appendix B. Specifically, Table 1 lists papers that are the first to audit a particular DP guarantee (e.g. [34, 82]), mechanism (e.g. [54, 72]), or achieve the current “best” auditing results (e.g., [82, 111]). Along with the relevant selected papers, we also report key details of the auditing methods introduced in these papers, which we extract through our systematization process as discussed later in Section 5.

Table 1 reveals that state-of-the-art work in DP auditing spans a wide range of mechanisms audited, from “simple” (e.g., Laplace, Noisy Histogram, etc.) to more complex machine learning and federated learning pipelines, vis-à-vis various DP guarantees and adjacency notions and employing different sub-sampling techniques. Thus, to support this wide range of mechanisms and techniques, auditing algorithms use different privacy attacks and evaluation functions, which we analyze in Sections 5.2 and 5.3, respectively. Overall, DP auditing techniques have achieved the three desiderata discussed in Section 3.3 only for simple mechanisms (e.g., Gaussian), while doing so for complex mechanisms remains an elusive task.

Although our selection was done in Jan 2025, we are confident it is meaningfully representative of the state of the art in DP auditing as of the time of submission; i.e., even if one were to add or remove a small number of papers to our list, our main takeaways would not substantially change.

**Out-of-scope Research.** As our systematization focuses on DP auditing, we did not include work estimating privacy leakage from algorithms using privacy metrics other than DP guarantees, e.g., using attack success rate and/or adversarial advantage [55, 92]. We also excluded research focusing on using static methods, such as type checkers, to verify the DP guarantees of an implementation without executing the program. While these methods are useful in various ways, they are arguably limited in the range of mechanisms they can audit, which makes them ineffective for verifying the DP guarantees of complex modern DP mechanisms, e.g., in the context of deep learning, as discussed in Section 3.1.

## 5 Systematizing DP Auditing Research Based on their Operational Modes

In this section, we systematize the state-of-the-art work on DP auditing based on the modes of operation specific to

Reference	Mechanisms Audited	Foundations		Operational Details		Progress		
		DP Guarantee	Sub-Sampling	Attack	Eval. Function	Efficient	End-to-End	Tight
Ding et al. [34]	Simple	$\epsilon$ -DP	–	–	Output-set	✗	✓	✓
Bischel et al. [18]	Simple	$\epsilon$ -DP	–	–	Output-set	✗	✓	✓
Jagielski et al. [54]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗	✓	✗
Nasr et al. [83]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗ ✗	✓ ✗	✗ ✓
Malek et al. [72]	ALIBI	Label DP	–	AI	Accuracy	✓	✓	✗
Houssiau et al. [52]	Synthetic Data	$(\epsilon, \delta)$ -DP	–	AI	FPR/FNR	✗	✓	✗
Askin et al. [12]	Simple	$\epsilon$ -DP	–	–	Distance Est.	✗	✓	✓
Lokna et al. [66]	Simple	$(\epsilon, \delta)$ -DP	–	DPD	Output-set	✗	✓	✓
Nasr et al. [82]	DP-SGD	$f$ -DP	Poisson –	MI	FPR/FNR	✗ ✓	✓ ✗	✗ ✓
Maddock et al. [70]	DP-FedSGD	User-level DP	Poisson	MI	FPR/FNR	✓	✗	✗
Pillutla et al. [91]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	Custom	✗ ✗	✓ ✗	✗ ✗
Steinke et al. [100]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	Accuracy	✓ ✓	✓ ✗	✗ ✗
Galen et al. [4]	DP-FedAvg	User-level DP	Shuffle	MI	Variance Est.	✓	✗	✗
Chadha et al. [21]	PATE	$(\alpha, \gamma)$ -RDP	–	MI	Custom	✗	✗	✗
Annamalai et al. [8]	Synthetic Data	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗ ✗	✓ ✗	✗ ✓
Feng et al. [45]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	–	Output-set	✗	✓	✓
Annamalai et al. [7]	DP-SGD	$\mu$ -GDP	–	MI	FPR/FNR	✗	✓	✗
Mahloujifar et al. [71]	DP-SGD Gaussian	$f$ -DP $\mu$ -GDP	Poisson –	Recon	Accuracy	✓ ✓	✗ ✓	✗ ✓
Annamalai et al. [6]	DP-SGD (Shuffle)	$(\epsilon, \delta)$ -DP	Shuffle	MI	FPR/FNR	✗	✗	✗
Cebere et al. [20]	DP-SGD	$\mu$ -GDP	Poisson	MI	FPR/FNR	✗	✓	✗
Xiang et al. [111]	DP-SGD Gaussian	$f$ -DP $\mu$ -GDP	Poisson –	MI	Accuracy	✓ ✓	✗ ✓	✗ ✓

**Table 1:** Summary of key prior work on DP auditing (see Appendix B for a full summary). NB: “Simple” mechanisms refer to Laplace, Report Noisy Max, Noisy Histogram, and Sparse Vector Technique mechanisms. Please see Sections 2, 3.3, and 5 for definitions of DP guarantees and sub-sampling methods, progress metrics, and attacks and evaluation functions, respectively.

DP auditing. Specifically, we extract the threat models used in the literature and introduce a new taxonomy to further break down these threat models into specific adversarial capabilities. Furthermore, we extract the types of privacy attacks and evaluation functions instantiated by auditing techniques. Finally, we discuss different techniques to calculate confidence levels for audits.

## 5.1 Threat Modeling

When auditing DP implementations, especially those deployed in the wild, it is important to do so vis-à-vis meaningful threat models. Naturally, these do not only affect the privacy guarantees of a mechanism but also impact whether the audit can be considered end-to-end as stronger adversarial models often require mechanisms to be modified before auditing. In this section, we focus on auditing

the DP-SGD algorithm [1] since the natural threat model for simple mechanisms like Laplace is already the *end-to-end* model.

In the context of private machine learning, different threat models have been used to instantiate privacy attacks (e.g., black-box, white-box, etc.). For more fine-grained analysis, we break them down into specific capabilities – namely, model visibility and types of canary used by the attacks – which enables us to systematize research along these axes independently.

**Model Visibility.** When training ML models using DP-SGD, only the final trained model is released and visible to the adversary.<sup>1</sup> However, the theoretical privacy analysis of DP-

<sup>1</sup>Except for specific settings, e.g., Federated Learning [75], where algorithms like DP-FedAvg and DP-FedSGD inherently disclose intermediate model updates.

SGD assumes that the intermediate models across all iterations are released as we do not know how to prove DP guarantees otherwise. Thus, early work assumes that the adversary has access to all intermediate models [82, 83]. More recently, the focus has shifted to auditing DP-SGD when only the final model output [7, 71, 91, 100] is available.

When only the final model is visible, the prototypical auditing algorithm outlined in Algorithm 1 should be modified so that the Challenger does not send all the intermediate outputs produced by the mechanism to the Adversary (line 8), but rather only the final output.

**Canaries.** DP definitions assume a strong adversary that can in theory define the entire dataset given as input to a given mechanism (i.e., “dataset canary”). However, in the context of DP-SGD, auditing in this threat model typically entails training on so-called “pathological” datasets, which can destroy model utility [83] or may not reflect real-world threat models. While these considerations do not affect the theoretical analysis of DP-SGD, which provides guarantees against the *worst-case* adversary, they can inform model developers of the privacy properties satisfied by the specific model being built.

Consequently, several DP-SGD auditing algorithms involve adversaries that can insert individual sample gradients at each step (i.e., “gradient canary”) [8, 20, 82]. In practice, however, adversaries may only be able to insert a single sample into the input dataset (i.e., “sample canary”). While early work has focused on inserting dataset canaries [83], more recent research relies on inserting sample canaries instead, which is much more restrictive [7, 71, 82, 91, 100].

To account for the use of canaries, in Algorithm 1, the Adversary defines the entire dataset (line 1) only in the dataset canary threat model. Whereas in the gradient and sample canary threat models, in line 1, the dataset  $D^-$  is (possibly randomly) sampled from a “natural” distribution, e.g., CIFAR-10, instead. Additionally, in the gradient canary threat model, the DP-SGD algorithm is modified to allow the Adversary to provide gradient canaries associated with the input canaries through the auxiliary information, which will be inserted by the algorithm in place of those computed on the input canaries.

## 5.2 Attacks

When auditing DP implementations of simpler algorithms (e.g., Laplace mechanism or Report Noisy Max) [18, 34, 65, 86], the algorithm can directly analyze the outputs of the mechanism. However, when auditing more complex algorithms (e.g., DP-SGD), it becomes crucial to rely on privacy attacks to post-process the raw outputs from mechanisms into observations that can be analyzed more easily. We discuss these attacks and their role next.

**Differential Privacy Distinguishability (DPD).** Here, we consider an adversary aiming to learn a “distinguishing function” to tell apart outputs from adjacent inputs [8, 10, 11, 17, 19, 66, 68, 118]. For instance, assume a machine learning

model is trained on outputs from adjacent datasets  $\mathcal{M}(D)$  and  $\mathcal{M}(D')$  for binary classification (‘0’ for  $\mathcal{M}(D)$  and ‘1’ for  $\mathcal{M}(D')$ ). Then, *fresh outputs* (not used in training) are sampled from the mechanism and converted into scalar “scores” by running inference on the model. The scalar scores represent the adversary’s confidence that the outputs came from processing  $D'$  as opposed to  $D$ .

**Membership Inference (MI).** When auditing differentially private machine learning algorithms (e.g., DP-SGD, DP synthetic data), a common post-processing technique is membership inference (MI) [5–8, 20, 21, 24, 31, 45, 47, 52, 54, 67, 70, 73, 82–84, 91, 100, 104, 110, 111, 114, 115] owing to its popularity as a privacy measurement tool for (non-private) ML models. In an MI attack, the adversary attempts to infer whether a given record was used to train the model. This strongly aligns with the add/remove adjacency where either a canary is added to the dataset or not.

Typically, when auditing in the Sample Canary threat model, MI is run by computing the model’s loss on the canary [5, 7, 8, 45, 54, 67, 71, 82–84, 91, 100, 104, 114, 115] or by training shadow models similar the target model (in terms of model/training architecture and data distribution) [8, 24, 31, 47, 52]. In the Gradient Canary threat model, the adversary can also compute the dot product [8, 70, 71, 82, 100, 100, 110, 111] or cosine similarity [4, 20, 73] between the gradient canary and noisy model updates or explicitly compute the likelihood ratio function [6].

**Attribute Inference (AI).** Although MI is the most commonly used privacy attack for auditing DP-ML techniques, in specific circumstances, such as Label DP [48] and synthetic data [8], Attribute Inference (AI) has also been used [52, 72]. In AI, the adversary attempts to infer specific private attributes of a given record, given the record’s public attributes. To perform auditing, Malek et al. [72] infer the label associated with the canary sample in Label DP, while Houssiau et al. [52] infer sensitive attributes of a tabular canary given all other attributes for DP synthetic data. Overall, establishing a baseline attack success rate is typically more complicated for AI [52].

**Reconstruction (Recon).** Drawing on the theoretical bounds by Hayes et al. [51], reconstruction attacks (Recon) can also be used to audit DP implementations [71]. Here, the adversary attempts to generate canaries from the data domain given only the output of the mechanism and the canary distribution. Although this can be difficult for arbitrary canary distributions, [71] shows that the problem is tractable when the canaries are uniformly sampled from discrete distributions each containing  $k$  canaries.

## 5.3 Evaluation Functions

Once simple scalar scores have been extracted from the mechanism outputs, there are different ways to fulfill the goal of the audit. While early work required the mechanism to be run multiple times ( $n_R > 1$ ), to reduce compu-

tational overhead, recent work has focused on doing so in a single run ( $n_R = 1$ ). In theory, the evaluation function can be specifically designed with the goal of the audit in mind—i.e., verification or estimation. However, in practice, evaluation functions are typically designed for the estimation goal, while verification is done by simply checking if the estimated parameters are consistent with the claimed DP guarantees—i.e., compute  $(\varepsilon_{\text{emp}}, \delta_{\text{emp}}) \leftarrow \zeta(\cdot)$  and check if  $\varepsilon_{\text{emp}} \leq \varepsilon$  and  $\delta_{\text{emp}} \leq \delta$  for verification. Next, we discuss estimation variants of different evaluation functions.

### 5.3.1 Multiple Runs

**Output-set optimization.** Early work on auditing DP implementations [18, 34, 65, 66, 86] focused on simple algorithms (e.g., Laplace mechanism or Report Noisy Max), directly testing the approximate DP condition with respect to a set of outputs  $E$ , i.e.,  $\Pr[\mathcal{M}(D) \in E] \leq e^\varepsilon \Pr[\mathcal{M}(D') \in E] + \delta$ . Since the outputs of these mechanisms are already simple scalars, the attack function  $\mathcal{P}$  simply returns the output of the mechanism without modifying it. To maximize the power of the audit, the evaluation function optimizes the output-set  $E$  that results in the largest  $\varepsilon$  for a given  $\delta$ .

**Distribution estimation.** Alternatively, one can empirically estimate the distributions the observations are drawn from. For instance, Askin et al. [12] and Gorla et al. [50] use kernel density estimators and histogram estimators, respectively, to estimate the distributions of observations drawn from  $\mathcal{M}(D)$  and  $\mathcal{M}(D')$ . After estimating the distribution, they test the approximate DP condition directly as done in output-set optimization. Conversely, Kong et al. [62] compute the Hockey-stick<sup>2</sup> or Rényi divergence depending on the claimed DP guarantees, while Huang et al. [53] estimate the variance of the distributions and estimate the privacy guarantees from the variance.

**FPR/FNR.** In the context of private ML, the most popular technique is to compute the false positive and false negative rates of the privacy attack and compare them with the bounds inferred from DP guarantees (cf. Section 2.4). First introduced by Jagielski et al. [54] for DP-SGD, this technique has been used in numerous research papers and extended to audit other DP mechanisms such as synthetic data generation, federated learning, and private prediction [5–8, 20, 21, 24, 31, 45, 47, 52, 67, 70, 73, 82–84, 104, 110, 114, 115].

One drawback is that this requires thousands, if not millions, of observations to provide reasonable confidence intervals [6]. Lu et al. [67] and Zanella-Beguelin et al. [115] attempt to reduce the number of runs required by optimizing the confidence interval computation but still require around a thousand runs to audit. Nasr et al. [82] reduce the number of runs to two by auditing each step of the DP-SGD algorithm independently. Similarly, Maddock et al. [70] only require a single run of the DP-FedSGD [76] algorithm, although they audit each round independently.

<sup>2</sup> $H_\alpha(P, Q) = \int_x \max\{P(x) - \alpha Q(x), 0\}$  [97].

**Multiple Canaries.** While previous work [34, 54, 82, 83] has focused on techniques that add only a single canary in each run, Pillutla et al. [91] use multiple canaries per run. Specifically, they perform multiple hypothesis tests (one for each canary) for each mechanism output, thus reducing the number of runs required to a few hundred. However, for computationally expensive algorithms like DP-SGD, even a few hundred runs can make it computationally prohibitive to audit in practice, especially for large language models [88].

### 5.3.2 One Run

**Guessing Game.** As mentioned, recent work has focused on reducing the number of runs needed to audit mechanisms to only one. This is formalized through a *guessing game* where DP guarantees are estimated using the *accuracy* of a privacy attack instead of FPR/FNRs. Malek et al. [72] introduce a heuristic to audit DP mechanisms by inserting multiple canaries into the mechanisms’ input. Steinke et al. [100] formalize this heuristic for MI and demonstrate that MI attacks targeting multiple canaries can indeed be used for DP auditing, albeit only achieving loose audits.

Mahloujifar et al. [71] and Xiang et al. [111] improve by auditing with, respectively,  $f$ -DP and information theory. Specifically, Mahloujifar et al. [71] extend [100] to cover  $f$ -DP, audit using Recon attacks, and improve the bound on attack success of Hayes et al. [51]. Xiang et al. [111] stick to MI but model the auditing task using information-theoretic principles (i.e., noisy channels).

Beyond fully connected neural networks and convolutional neural networks, recent work by Panda et al. [88] use the techniques from [100] to optimize canary design and audit large language models.

**Variance estimation.** In the context of Federated Learning, Galen et al. [4] audit the DP-FedAvg algorithm by inserting multiple independent isotropically distributed canaries into each round of the algorithm. They then estimate the mean and variance of the cosine similarities between the inserted canaries and the mechanism’s output.

## 5.4 Confidence Level

Whether the audit goal is verification or estimation, as mentioned in Section 3.1, performing it at a given confidence level can be important to provide auditors with statistical assurance that the result of the audit is correct.

One popular method is to estimate the type I and type II errors of the privacy attack to a given level of confidence using statistical tools like the Clopper-Pearson [30] interval [5–8, 10, 19–21, 24, 47, 52, 54, 66, 70, 72, 73, 82–84, 86, 104, 114, 115], Katz-log [60] interval [67, 82], or credible intervals [4, 31, 82, 110, 115]. Some audits also use custom methods to specify the confidence level tuned to their specific auditing algorithms [11, 12, 17, 19, 34, 50, 62, 71, 91, 100, 111]. However, some do not report any confidence intervals at all [18, 53, 65, 68, 107, 118].

## 5.5 Remarks

In this section, we reviewed and extracted key details of auditing techniques, including the types of mechanisms and DP guarantees audited, privacy attacks and evaluation functions employed, and the progress made along the three desiderata discussed in Section 3.3 (i.e., efficiency, end-to-end-ness, and tightness).

While DP audits have used a wide range of techniques, very few techniques have achieved all three desired properties, and only for simple DP mechanisms. Next, we discuss the main research findings based on our systematic review of the DP auditing body of work.

## 6 Towards Efficient, End-to-End, and Tight DP Audits

We now discuss several key details of DP auditing that have been neglected by prior work, discover limiting factors to achieving all three desiderata, and identify novel directions for future work to focus on.

### 6.1 DP Guarantees & Adjacency

Naturally, the DP guarantee variant being audited is a fundamental requirement of any DP auditing algorithm, as per Definition 3.2. As evident from Table 1, while earlier work [12, 19, 34, 52, 54, 66, 83] focuses on simpler variants of DP such as  $\epsilon$ -DP and  $(\epsilon, \delta)$ -DP, more recent audits [7, 20, 71, 82, 111] do so on  $f$ -DP and  $\mu$ -GDP instead to achieve tight audits of mechanisms.

This is because simpler variants of DP may not always describe the privacy leakage of specific mechanisms optimally, i.e., the *actual* privacy leakage of the mechanism may be far from the *theoretical* privacy leakage expected from simple DP guarantees. In turn, audits are loose since DP auditing measures the actual privacy leakage from mechanisms. For instance, Nasr et al. [82] achieve tight audits of DP-SGD using the  $f$ -DP guarantee, which they show models the actual privacy leakage of DP-SGD much more closely compared to  $(\epsilon, \delta)$ -DP initially used for auditing [54, 83]. Therefore, recent audits of DP-SGD use the  $f$ -DP or the related  $\mu$ -GDP guarantee instead [7, 20, 71, 82, 111].

On the other hand, auditing with  $f$ -DP poses several challenges. Firstly, as discussed in Section 6.2,  $f$ -DP guarantees of mechanisms are not always known (this is an active area of research). Furthermore, when used for DP estimation, it can be difficult to estimate the full  $f$ -DP curve compared to DP guarantees with only a few scalar parameters (e.g.,  $\epsilon$ -DP,  $(\epsilon, \delta)$ -DP,  $\mu$ -GDP). In fact, even when estimating  $(\epsilon, \delta)$ -DP, the  $\delta$  parameter is typically fixed [6, 8, 12, 45, 52, 54, 66, 83, 91, 100] so that only  $\epsilon$  is being estimated. Therefore, there are no known methods to estimate  $f$ -DP guarantees with confidence intervals, and doing so constitutes an open research question.

Moreover, the DP guarantees and adjacency notions audited might not always be explicitly stated. This is the case

for the adjacency notion in several key papers [7, 11, 45, 53, 84]. Furthermore, several results [7, 20, 71, 82, 111] are given with respect to  $(\epsilon, \delta)$ -DP, although closer inspection of their methodology reveals that they are instead auditing the  $f$ -DP or  $\mu$ -GDP guarantees.

#### Takeaway 1

DP audits should be clear about the exact DP guarantee being audited along with the adjacency notions considered. In general, auditing  $f$ -DP guarantees are optimal, but deriving tight guarantees for specific mechanisms is currently an open challenge.

### 6.2 Sub-sampling Schemes

Another key detail when auditing mechanisms is the underlying sub-sampling scheme. From Table 1, we observe that the most common sub-sampling scheme audited is Poisson [8, 20, 45, 54, 70, 71, 82, 83, 91, 100, 111], owing to its use in the DP-SGD algorithm. However, recent work [4, 6] has also looked at auditing the shuffling scheme. Overall, while audits not involving sub-sampling schemes are known to be tight [12, 19, 34, 66, 71, 82, 111], tight audits involving sub-sampling schemes are much rarer [8, 45, 83], usually requiring additional assumptions, such as step-by-step auditing or simple models.

In this context, the main challenge is having tight DP guarantees of mechanisms employing sub-sampling schemes, i.e., knowing that the *actual* privacy leakage of mechanisms matches the *theoretical* privacy leakage expected from the DP guarantee. Next, we focus on the Gaussian Mechanism since sub-sampling schemes are most commonly used with that, e.g., in DP-SGD [1]. In Table 2, we summarize the sub-sampling schemes and adjacency notions for which tight DP guarantees are currently known. While tight guarantees are known for Poisson sub-sampling, Sampling WR, and Balls-in-Bins sampling for the Add/Remove, Edit, and Zero-out adjacencies, respectively, for the vast majority of sub-sampling scheme and adjacency pairs, tight guarantees remain unknown.

#### Takeaway 2

Deriving the tight guarantees for many popular sub-sampling schemes is an open research question; thus, the ability to tightly audit mechanisms that use them is currently limited.

### 6.3 Threat Models

One elusive combination of desired properties for audits of modern complex mechanisms like DP-SGD is for them to be both end-to-end and tight. Specifically, Table 1 reveals that end-to-end audits of DP-SGD are not tight [7, 20, 54, 71, 82, 83, 91, 100] and viceversa [82, 83]. The only exception is the work by Feng et al. [45], who achieve both tight and

		Sub-Sampling Scheme				
		Poisson	Sampling WR	Sampling WOR	Shuffling	Balls-in-Bins
Adjacency	Add/Remove	[119]		[13]		
	Edit		[98]	[119]		
	Zero-Out	[28]			[28, 29]	[27]

**Table 2:** Summary of whether tight guarantees of Gaussian Mechanism are currently known for sub-sampling schemes and adjacency notions. NB: Cell colored green if prior work has identified tight guarantees, yellow if there have been prior work but tight guarantees are not currently known, and brown otherwise.

		Canary Type		
		Sample	Gradient	Dataset
Model Visibility	Intermediate Models	[53]	[17, 71, 73, 82, 83, 91, 100, 111]	[83]
	Final Model	[31, 45, 54, 67, 71, 82–84, 88, 100, 111, 115]	[20]	[5, 11, 83]

**Table 3:** Summary of different threat models for DP-SGD auditing considered in prior work. NB: Cell colored green if tight auditing has been achieved and brown otherwise.

end-to-end audits of DP-SGD, but only for simple logistic regression models.

To analyze this further, we break down end-to-end audits into specific adversarial capabilities (model visibility and canary type), as explained in Section 5.1. In Table 3, we highlight the threat models where state-of-the-art DP-SGD audits are tight. Generally, tight auditing in stronger models has been accomplished; however, while there has been significant work on the weakest (Final Model, Sample Canary) model, aka “Black-box,” less work has focused on the slightly stronger (Final Model, Gradient Canary) and (Intermediate Models, Sample Canary) models.<sup>3</sup>

Overall, we consider a DP audit “end-to-end” only in the weakest threat model (Final Model, Sample Canary). Given that these audits are still far from tight, focusing efforts on the two slightly stronger threat models would be promising interim directions.

#### Takeaway 3

More research on auditing DP-SGD in the (Final Model, Gradient Canary) and (Intermediate Models, Sample Canary) threat models can potentially pave the way toward tight audits in (Final Model, Sample Canary), which is crucial for end-to-end audits.

only require two runs but they only manage to audit a single step of the DP-SGD mechanism. Steinke et al. [100], Mahloujifar et al. [71], Panda et al. [88], and Xiang et al. [111] all propose techniques to audit the entire DP-SGD mechanism in one run. However, even in the strong (Gradient Canary, Intermediate Models) threat model, their audits remain loose, at best only achieving an empirical lower bound of  $\epsilon_{\text{emp}} \approx 4.5$  for a theoretical  $\epsilon = 8$  [111]. Similarly, one-run audits of Federated Learning variants of DP-SGD (i.e., DP-FedAvg and DP-FedSGD) are also loose [4, 70]. By contrast, tight one-run audits for simpler mechanisms such as the Gaussian Mechanism are possible [111]. This suggests there may be fundamental challenges that may need to be overcome when auditing DP-SGD in one run—we discuss this in more detail in Section 7.4.

#### Takeaway 4

Efficient auditing in a single run is a primary focus of recent DP auditing work but, even in powerful threat models, e.g., (Gradient Canary, Intermediate Models), audits remain loose. Given the complications surrounding auditing in weaker threat models (see Section 6.3), research should focus on achieving tight audits in one run in these more powerful threat models first.

## 6.4 Evaluation Functions

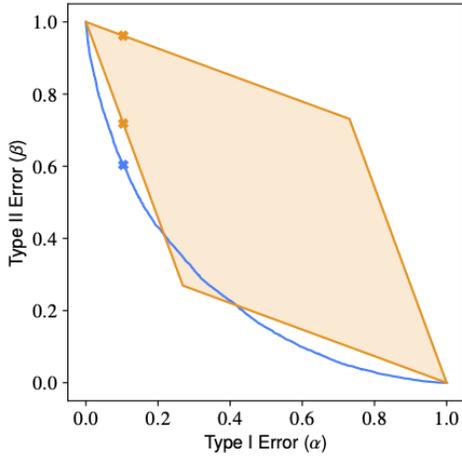
Another elusive combination of the desiderata is accomplishing DP audits that are both efficient and tight. From Table 1, we observe that only recently this has been achieved, albeit only for simple mechanisms like Gaussian [71, 111]. As a result, an important problem for future research is designing evaluation functions that only require mechanisms to be run once for auditing.

As discussed in Section 5.3, Nasr et al. [82]’s techniques

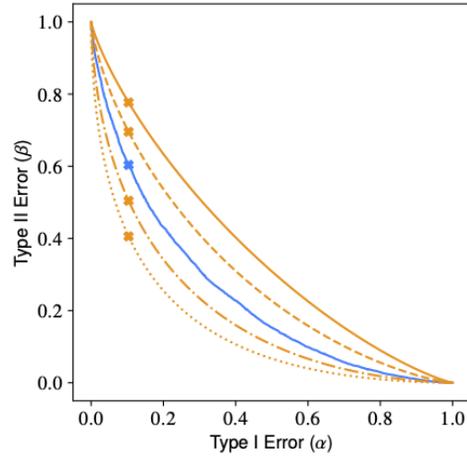
<sup>3</sup>To ease presentation, we use the (Model Visibility, Canary Type) notation to quickly refer to the corresponding threat models.

## 6.5 Privacy Region & Confidence Level

The concept of *privacy region* has been used by auditors to both verify and estimate DP guarantees of the mechanisms under study. In Figure 2, we plot how the Gaussian mechanism ( $\sigma = 1$ ) gets audited under both the  $(\epsilon, \delta)$ -DP and  $\mu$ -GDP variants. More precisely, in Figure 2a, the privacy region is estimated from the mechanism and compared against the hypothetical guarantees of  $(1, 10^{-5})$ -DP, which the mechanism does *not* satisfy. Note the presence of large sections of the estimated privacy region that fall well outside the shaded theoretical privacy region for  $(1, 10^{-5})$ -DP. This suggests that the claimed DP guarantees are wrong and they



(a) Verifying (incorrectly) claimed  $(1, 10^{-5})$ -DP guarantees.



(b) Estimating empirical privacy guarantees using  $\mu$ -GDP, with  $\mu = \{0.5, 0.75, 1.25, 1.5\}$ .

**Figure 2:** Privacy regions when auditing Gaussian mechanism ( $\sigma = 1$ ) with  $(\epsilon, \delta)$ -DP and  $\mu$ -GDP. The blue line indicates estimated guarantees and the orange lines (and shaded regions) represent claimed or candidate guarantees.

can therefore be refuted experimentally.

On the other hand, in Figure 2b, the estimated privacy region is compared against multiple “candidate”  $\mu$ -GDP regions ( $\mu = 0.5, 0.75, 1.25, 1.5$ ) to estimate the empirical  $\mu_{\text{emp}}$ -GDP guarantees of the mechanism. The  $\mu$  parameter is changed to plot different theoretical privacy regions (for simplicity, we only plot the lower boundaries). By doing so, we can estimate that the mechanism satisfies  $\mu_{\text{emp}}$ -GDP for  $0.75 \leq \mu_{\text{emp}} \leq 1.25$ , which we can further fine-tune by comparing  $\mu$ -GDP regions at lower resolutions.

In both cases, either the full privacy region can be estimated (represented by the line), or a single point estimate can be used (represented by the markers). This choice mainly affects the *confidence level* of the overall audit procedure as confidence levels derived for single-point estimates may not necessarily transfer to full-region estimates.

DP auditing research typically bypasses computing valid confidence intervals to ease experimentation. As mentioned in Section 5.4, some researchers either leave that out altogether [18, 53, 65, 68, 107, 118] or they estimate the full privacy region but report confidence intervals for a single-point estimate instead [7, 70, 82, 83], thus making them invalid. Nevertheless, in practice, audits should ensure that valid confidence intervals are computed for the auditing method used. For instance, the best possible single-point estimate can be derived using “shadow models” [8] or the entire region can be estimated using  $\mu$ -GDP [82]. Additionally, further research could also focus on reliably estimating full privacy regions with meaningful confidence intervals.

#### Takeaway 5

When deploying auditing algorithms in the real world, care has to be taken to ensure that the confidence intervals computed are *valid*. Further research could also focus on doing so when estimating full privacy regions.

## 7 Discussion & Conclusion

This paper systematized research on DP auditing using a systematic framework geared to categorize the underlying settings/techniques and measure progress in the state of the art along three key desired properties: *efficiency*, *end-to-end-ness*, and *tightness*.

Overall, our analysis showed that while DP auditing has made significant and rapid progress in making audits more effective, useful, and easily deployable against production systems, there are still areas with substantial research gaps. In the process, we identified several key insights and open problems; we are confident that future work can build upon these to establish a “Gold Standard” for DP auditing.

In the rest of this section, we review related work, present actionable recommendations, and highlight open research questions future work should prioritize.

### 7.1 Related Work

To the best of our knowledge, our work is the first to systematically review and analyze the body of work on DP auditing. Namatevs et al. [80] recently present a survey of DP auditing research; however, they do not analyze relevant work within a systematic framework. By contrast, we do so to comprehensively analyze the state of the art, measure progress, identify gaps, and highlight open research problems primarily along three key characteristics: whether audits are efficient, end-to-end, and tight.

Systematization of knowledge efforts in related areas include [90], who present attacks and defenses on (non-private) machine learning systems using a comprehensive adversarial framework. Salem et al. [96] focus on privacy risks and provide a unifying framework relying on privacy games. Finally, Desfontaines and Pejó [33] systematize the various variants and extensions of DP. While DP auditing relies on and is informed by the notions of privacy attacks,

DP Guarantee	End-to-End	Efficient	Not Efficient
$(\epsilon, \delta)$ -DP, $\epsilon$ -DP	✓ ✓	Xiang et al. [111]	Jagielski et al. [54], Lokna et al. [66]
$(\alpha, \gamma)$ -RDP	✗	–	Chadha et al. [21]
$f$ -DP, $\mu$ -GDP	✗ ✓	Nasr et al. [82] Xiang et al. [111]	Cebere et al. [20] Annamalai et al. [7]

**Table 4:** DP auditing methods considered to be state-of-the-art (SOTA) with respect to tightness, broken down by DP guarantee, efficiency, and end-to-end-ness. (To ease presentation, we omit *et al.* and only report the last name of each paper’s first author.)

games, and DP variants, our work focuses on the specific process of deriving empirical bounds, which has thus far not been explored in prior work.

## 7.2 What is considered SOTA in DP Auditing?

Although there are different settings and thus different auditing techniques in literature, we attempt to identify the results that can be considered state-of-the-art (SOTA) with respect to tightness. More precisely, in Table 4, we do so while breaking down the DP guarantees audited and whether the other two key desirable properties are achieved.

Put simply, when efficient audits are required, Xiang et al. [110]’s work is SOTA. Specifically, their techniques achieve tight audits for simple mechanisms such as Laplace, Gaussian, or Randomized Response but only loose audits for more complex mechanisms like DP-SGD. One main limitation of Xiang et al. [111]’s work is the potential loss in model utility as it requires millions of target records to be added to the initial dataset for audits to be tight. For DP-SGD specifically, loosening the end-to-end property allows us to achieve tight audits using Nasr et al. [82]’s technique, while also preserving model utility.

When audits can be inefficient but minimal utility loss is expected, we consider the auditing method by Jagielski et al. [54] to be SOTA for complex algorithms like DP-SGD and Lokna et al. [66] for simple mechanisms satisfying  $(\epsilon, \delta)$ -DP. We also note that Chadha et al. [21]’s method is the *only* technique available for mechanisms satisfying  $(\alpha, \gamma)$ -RDP. Finally, for  $f$ -DP (or  $\mu$ -GDP), Annamalai et al. [7]’s auditing method is SOTA for end-to-end audits, while that by Cebere et al. [20] is SOTA for audits over sub-sampling that do not need to be end-to-end.

## 7.3 Recommendations

**Auditing Details.** As discussed in Section 6.1, one not-so-uncommon limitation in DP auditing research is that key details of the auditing methodology are spread across various sections of the paper and may not be explicitly stated or left unclear. We suggest that researchers working on DP auditing explicitly state the DP guarantee audited, along with the adjacency notion, sub-sampling scheme, threat model, and confidence interval considered.

## Audit Card

**Mechanism:** DP-SGD  
**(Claimed) DP Guarantee:**  $(1.0, 10^{-5})$ -DP  
**Adjacency:** Add/Remove  
**Sub-sampling:** Poisson  
**Threat Model:** (Intermediate Models, Gradient Canary)  
**No. of Runs:** 1K  
**Confidence level:** 95%  
**Audit Result:**  $(0.95, 10^{-5})$ -DP (not violated)

**Figure 3:** Example of a Possible Design for Audit Cards.

For instance, they could adopt an “Auditing Card” to be added in the appendix and/or in code repositories, similar to the “Model Cards” released for models published on HuggingFace.<sup>4</sup> We include a draft template in Figure 3; future work could conduct a user study with researchers and practitioners working on DP auditing and explore the viability and opportunity of this approach.

**Tighter  $f$ -DP or  $\mu$ -GDP guarantees.** As also discussed by Gomez et al. [49], we also advocate for the standardized release and auditing of tighter  $f$ -DP or  $\mu$ -GDP guarantees. For example, when auditing the Gaussian mechanism, future audits could estimate and compare the empirical  $\mu_{\text{emp}}$  value to the theoretical  $\mu$ -GDP guarantee, instead of converting both to a  $(\epsilon, \delta)$ -DP value with an arbitrarily chosen  $\delta$ . For mechanisms without a  $\mu$ -GDP guarantee, the full  $f$ -DP curve could be estimated, and the  $\Delta^{\leftarrow}$ -divergence [49, 59] between the estimated and theoretical curve could be used as a “goodness of fit” metric.

## 7.4 Open Research Questions

In conclusion, our systematization effort highlights the presence of at least three open research questions.

**RQ1: Can we derive tight guarantees for all sub-sampling schemes?** Although tight theoretical guarantees have been derived for many popular sub-sampling schemes (e.g., Poisson), much less is known about schemes like sampling without replacement and shuffling. These schemes are commonly used to train models efficiently [6], and in the context of auditing, having tight guarantees is necessary for audits to be tight. Thus, deriving tight theoretical guarantees for these schemes remains an important open problem.

**RQ2: Are end-to-end and tight audits possible?** One of the most elusive combinations of desired properties is end-to-end-ness and tightness. While this is possible for simple mechanisms, for complex algorithms like DP-SGD, very little work has come close. Based on our systematic analysis, we believe this question can be further broken down into investigating slightly weaker threat models, namely, (Final Model, Gradient Canary) and (Intermediate

<sup>4</sup>See <https://huggingface.co/docs/hub/en/model-cards>.

Models, Sample Canary) before addressing the end-to-end threat model, i.e., (Final Model, Sample Canary). Ceberé et al. [20] audit DP-SGD in (Final Model, Gradient Canary) but fail to achieve tightness, suggesting that tighter theoretical analysis might be necessary, which is in itself an open research question [5, 84]. Hence, we call for future work to focus on the (Final Model, Gradient Canary) and (Intermediate Models, Sample Canary) threat models.

### RQ3: Are efficient and tight audits of DP-SGD possible?

Another elusive combination of properties is efficiency and tightness. While several methods have been suggested toward this goal [71, 100, 111], only recently Xiang et al. [111] present efficient and tight audits, albeit only for simple mechanisms. However, while recent work has identified fundamental challenges in efficiently auditing complex mechanisms [61], these challenges do not particularly limit DP-SGD; thus, efficient and tight audits of DP-SGD could be possible. Similar to RQ2, it might be useful to first consider alternative threat models to elicit the maximal privacy leakage in efficient audits.

## References

- [1] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep Learning with Differential Privacy. In *CCS*, 2016.
- [2] J. M. Abowd, R. Ashmead, R. Cumings-Menon, S. Garfinkel, M. Heineck, C. Heiss, R. Johns, D. Kifer, P. Leclerc, A. Machanavajjhala, et al. The 2020 census disclosure avoidance system topdown algorithm. *Harvard Data Science Review*, 2022.
- [3] A. Albarghouthi and J. Hsu. Synthesizing coupling proofs of differential privacy. In *POPL*, 2017.
- [4] G. Andrew, P. Kairouz, S. Oh, A. Oprea, H. B. McMahan, and V. Suriyakumar. One-shot Empirical Privacy Estimation for Federated Learning. In *ICLR*, 2024.
- [5] M. S. M. S. Annamalai. It’s Our Loss: No Privacy Amplification for Hidden State DP-SGD With Non-Convex Loss. In *AISec*, 2024.
- [6] M. S. M. S. Annamalai, B. Balle, E. De Cristofaro, and J. Hayes. To Shuffle or not to Shuffle: Auditing DP-SGD with Shuffling. *arXiv:2411.10614*, 2024.
- [7] M. S. M. S. Annamalai and E. De Cristofaro. Nearly Tight Black-Box Auditing of Differentially Private Machine Learning. In *NeurIPS*, 2024.
- [8] M. S. M. S. Annamalai, G. Ganey, and E. De Cristofaro. “What do you want from theory alone?” Experimenting with Tight Auditing of Differentially Private Synthetic Data Generation. In *USENIX Security*, 2024.
- [9] Apple Differential Privacy Team. Learning with Privacy at Scale. <http://docs-assets.developer.apple.com/ml-research/papers/learning-with-privacy-at-scale.pdf>, 2017.
- [10] H. H. Arcolezi and S. Gambs. Revealing the True Cost of Locally Differentially Private Protocols: An Auditing Perspective. *PoPETS*, 2024.
- [11] Ö. Askin, H. Dette, M. Dunsche, T. Kutta, Y. Lu, Y. Wei, and V. Zikas. General-Purpose  $f$ -DP Estimation and Auditing in a Black-Box Setting. *arXiv:2502.07066*, 2025.
- [12] Ö. Askin, T. Kutta, and H. Dette. Statistical Quantification of Differential Privacy: A Local Approach. In *IEEE S&P*, 2022.
- [13] B. Balle, G. Barthe, and M. Gaboardi. Privacy Amplification by Subsampling: Tight Analyses via Couplings and Divergences. In *NeurIPS*, 2018.
- [14] G. Barthe, G. Danezis, B. Grégoire, C. Kunz, and S. Zanella-Beguelin. Verified Computational Differential Privacy with Applications to Smart Metering. In *CSF*, 2013.
- [15] G. Barthe, M. Gaboardi, E. J. Gallego Arias, J. Hsu, A. Roth, and P.-Y. Strub. Higher-Order Approximate Relational Refinement Types for Mechanism Design and Differential Privacy. *SIGPLAN Notices*, 2015.
- [16] G. Barthe, B. Köpf, F. Olmedo, and S. Zanella Beguelin. Probabilistic relational reasoning for differential privacy. In *POPL*, 2012.
- [17] D. Bernau, G. Eibl, P. Grassal, H. Keller, and F. Kerschbaum. Quantifying identifiability to choose and audit  $\epsilon$  in differentially private deep learning. In *VLDB*, 2021.
- [18] B. Bichsel, T. Gehr, D. Drachler-Cohen, P. Tsankov, and M. Vechev. DP-Finder: Finding Differential Privacy Violations by Sampling and Optimization. In *CCS*, 2018.
- [19] B. Bichsel, S. Steffen, I. Bogunovic, and M. Vechev. DP-Sniper: Black-Box Discovery of Differential Privacy Violations using Classifiers. In *IEEE S&P*, 2021.
- [20] T. Ceberé, A. Bellet, and N. Papernot. Tighter Privacy Auditing of DP-SGD in the Hidden State Threat Model. In *ICLR*, 2025.
- [21] K. Chadha, M. Jagieski, N. Papernot, C. A. Choquette-Choo, and M. Nasr. Auditing Private Prediction. In *ICML*, 2024.
- [22] K. Chaudhuri and N. Mishra. When Random Sampling Preserves Privacy. In *Crypto*, 2006.
- [23] R. Chen, Q. Xiao, Y. Zhang, and J. Xu. Differentially Private High-Dimensional Data Publication via Sampling-Based Inference. In *KDD*, 2015.
- [24] K. Chida, S. Kakuta, H. Itakura, I. Ishihara, K. Yoshioka, and H. Takeuchi. Experimental Evaluation for Risk Assessment of Privacy Preserving Synthetic Data. In *ICMDAI*, 2024.
- [25] C. A. Choquette-Choo, A. Ganesh, S. Haque, T. Steinke, and A. G. Thakurta. Near-Exact Privacy Amplification for Matrix Mechanisms. In *ICLR*, 2025.
- [26] R. Chourasia, J. Ye, and R. Shokri. Differential Privacy Dynamics of Langevin Diffusion and Noisy Gradient Descent. In *NeurIPS*, 2021.
- [27] L. Chua, B. Ghazi, C. Harrison, P. Kamath, R. Kumar, E. J. Leeman, P. Manurangsi, A. Sinha, and C. Zhang. Balls-and-Bins Sampling for DP-SGD. In *AISTATS*, 2025.

- [28] L. Chua, B. Ghazi, P. Kamath, R. Kumar, P. Manurangsi, A. Sinha, and C. Zhang. How Private are DP-SGD Implementations? In *ICML*, 2024.
- [29] L. Chua, B. Ghazi, P. Kamath, R. Kumar, P. Manurangsi, A. Sinha, and C. Zhang. Scalable DP-SGD: Shuffling vs. Poisson Subsampling. *NeurIPS*, 2024.
- [30] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 1934.
- [31] E. Debenedetti, G. Severi, N. Carlini, C. A. Choquette-Choo, M. Jagielski, M. Nasr, E. Wallace, and F. Tramèr. Privacy Side Channels in Machine Learning Systems. In *USENIX Security*, 2024.
- [32] D. Desfontaines. A list of real-world uses of differential privacy. <https://desfontain.es/blog/real-world-differential-privacy.html>, 2025.
- [33] D. Desfontaines and B. Pejó. SoK: Differential privacies. *PoPETS*, 2020.
- [34] Z. Ding, Y. Wang, G. Wang, D. Zhang, and D. Kifer. Detecting Violations of Differential Privacy. In *CCS*, 2018.
- [35] K. Dixit, M. Jha, S. Raskhodnikova, and A. Thakurta. Testing the Lipschitz Property over Product Distributions with Applications to Data Privacy. In *TCC*, 2013.
- [36] J. Dong, A. Roth, and W. J. Su. Gaussian Differential Privacy. *arXiv:1905.02383*, 2019.
- [37] C. Dwork. Differential Privacy. In *ICALP*, 2006.
- [38] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *EUROCRYPT*, 2006.
- [39] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC*, 2006.
- [40] C. Dwork, M. Naor, O. Reingold, G. N. Rothblum, and S. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *STOC*, 2009.
- [41] C. Dwork, A. Roth, et al. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science*, 2014.
- [42] Elsevier. Scopus: Comprehensive, multidisciplinary, trusted abstract and citation database. <https://www.elsevier.com/products/scopus>, 2025.
- [43] Ú. Erlingsson, V. Pihur, and A. Korolova. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In *CCS*, 2014.
- [44] V. Feldman, I. Mironov, K. Talwar, and A. Thakurta. Privacy Amplification by Iteration. In *FOCS*, 2018.
- [45] S. Feng and F. Tramèr. Privacy Backdoors: Stealing Data with Corrupted Pretrained Models. In *ICML*, 2024.
- [46] M. Gaboardi, A. Haeberlen, J. Hsu, A. Narayan, and B. C. Pierce. Linear dependent types for differential privacy. In *POPL*, 2013.
- [47] G. Ganev, M. S. M. S. Annamalai, and E. De Cristofaro. The Elusive Pursuit of Replicating PATE-GAN: Benchmarking, Auditing, Debugging. *TMLR*, 2025.
- [48] B. Ghazi, N. Golowich, R. Kumar, P. Manurangsi, and C. Zhang. Deep Learning with Label Differential Privacy. In *NeurIPS*, 2021.
- [49] J. F. Gomez, B. Kulynych, G. Kaissis, J. Hayes, B. Balle, and A. Honkela.  $(\epsilon, \delta)$  Considered Harmful: Best Practices for Reporting Differential Privacy Guarantees. *arXiv:2503.10945*, 2025.
- [50] D. Gorla, L. Jalouzet, F. Granese, C. Palamidessi, and P. Piantanida. On the (im)possibility of estimating various notions of differential privacy. In *ICTCS*, 2023.
- [51] J. Hayes, B. Balle, and S. Mahloujifar. Bounding Training Data Reconstruction in DP-SGD. In *NeurIPS*, 2023.
- [52] F. Houssiau, J. Jordon, S. N. Cohen, A. Elliott, J. Geddes, C. Mole, C. Rangel-Smith, and L. Szpruch. TAPAS: a Toolbox for Adversarial Privacy Auditing of Synthetic Data. In *SyntheticData4ML Workshop NeurIPS*, 2022.
- [53] W. Huang, Z. Zhang, W. Zhao, J. Peng, W. Xu, Y. Liao, S. Zhou, and Z. Wang. Auditing privacy budget of differentially private neural network models. *Neurocomputing*, 2025.
- [54] M. Jagielski, J. Ullman, and A. Oprea. Auditing Differentially Private Machine Learning: How Private is Private SGD? In *NeurIPS*, 2020.
- [55] B. Jayaraman and D. Evans. Evaluating Differentially Private Machine Learning in Practice. In *USENIX Security*, 2019.
- [56] M. Johnson. GitHub Pull Request: fix prng key reuse in differential privacy example. <https://github.com/google/jax/pull/3646>, 2020.
- [57] P. Kairouz, B. McMahan, S. Song, O. Thakkar, A. Thakurta, and Z. Xu. Practical and Private (Deep) Learning Without Sampling or Shuffling. In *ICML*, 2021.
- [58] P. Kairouz, S. Oh, and P. Viswanath. The Composition Theorem for Differential Privacy. In *ICML*, 2015.
- [59] G. Kaissis, S. Kolek, B. Balle, J. Hayes, and D. Rueckert. Beyond the Calibration Point: Mechanism Comparison in Differential Privacy. In *ICML*, 2024.
- [60] D. Katz, J. Baptista, S. Azen, and M. Pike. Obtaining Confidence Intervals for the Risk Ratio in Cohort Studies. *Biometrics*, 1978.
- [61] A. Keinan, M. Shenfeld, and K. Ligett. How Well Can Differential Privacy Be Audited in One Run? *arXiv:2503.07199*, 2025.
- [62] W. Kong, A. M. Medina, M. Ribero, and U. Syed. DP-Auditorium: A Large Scale Library for Auditing Differential Privacy. In *IEEE S&P*, 2024.
- [63] T. Labs. Safely use, share & sell insights from personal data. <https://www.tmlt.io/>, 2025.
- [64] J. Lee and C. W. Clifton. Top-k frequent itemsets via differentially private fp-trees. In *KDD*, 2014.
- [65] X. Liu and S. Oh. Minimax Optimal Estimation of Approximate Differential Privacy on Neighboring Databases. In *NeurIPS*, 2019.
- [66] J. Lokna, A. Paradis, D. I. Dimitrov, and M. Vechev. Group and Attack: Auditing Differential Privacy. In *CCS*, 2023.

- [67] F. Lu, J. Munoz, M. Fuchs, T. LeBlond, E. Zaresky-Williams, E. Raff, F. Ferraro, and B. Testa. A General Framework for Auditing Differentially Private Machine Learning. In *NeurIPS*, 2022.
- [68] Y. Lu, M. Magdon-Ismael, Y. Wei, and V. Zikas. Eureka: A General Framework for Black-box Differential Privacy Estimators. In *IEEE S&P*, 2024.
- [69] M. Lyu, D. Su, and N. Li. Understanding the Sparse Vector Technique for Differential Privacy. *VLDB*, 2017.
- [70] S. Maddock, A. Sablayrolles, and P. Stock. CANIFE: Crafting Canaries for Empirical Privacy Measurement in Federated Learning. In *ICLR*, 2023.
- [71] S. Mahloujifar, L. Melis, and K. Chaudhuri. Auditing  $f$ -Differential Privacy in One Run. *arXiv:2410.22235*, 2024.
- [72] M. Malek Esmaeili, I. Mironov, K. Prasad, I. Shilov, and F. Tramer. Antipodes of Label Differential Privacy: PATE and ALIBI. In *NeurIPS*, 2021.
- [73] M. Matsumoto, T. Takahashi, S. P. Liew, and M. Oguchi. Measuring Local and Shuffled Privacy of Gradient Randomized Response. *Journal of Information Processing*, 2024.
- [74] B. McMahan and A. Thakurta. Federated Learning with Formal Differential Privacy Guarantees. <https://research.google/blog/federated-learning-with-formal-differential-privacy-guarantees/>, 2022.
- [75] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang. Learning differentially private recurrent language models. *arXiv:1710.06963*, 2017.
- [76] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang. Learning Differentially Private Recurrent Language Models. In *ICLR*, 2018.
- [77] Microsoft. IOM and Microsoft release first-ever differentially private synthetic dataset to counter human trafficking. <https://www.microsoft.com/en-us/research/blog/iom-and-microsoft-release-first-ever-differentially-private-synthetic-dataset-to-counter-human-trafficking/>, 2022.
- [78] I. Mironov. On significance of the least significant bits for differential privacy. In *CCS*, 2012.
- [79] I. Mironov. Rényi Differential Privacy. In *CSF*, 2017.
- [80] I. Namatevs, K. Sudars, A. Nikulins, and K. Ozols. Privacy Auditing in Differential Private Machine Learning: The Current Trends. *MDPI Applied Sciences*, 2025.
- [81] A. Narayanan and V. Shmatikov. Robust De-anonymization of Large Sparse Datasets. In *IEEE S&P*, 2008.
- [82] M. Nasr, J. Hayes, T. Steinke, B. Balle, F. Tramèr, M. Jagielski, N. Carlini, and A. Terzis. Tight Auditing of Differentially Private Machine Learning. In *USENIX Security*, 2023.
- [83] M. Nasr, S. Song, A. Thakurta, N. Papernot, and N. Carlini. Adversary Instantiation: Lower Bounds for Differentially Private Machine Learning. In *IEEE S&P*, 2021.
- [84] M. Nasr, T. Steinke, A. Ganesh, B. Balle, C. A. Choquette-Choo, M. Jagielski, J. Hayes, A. G. Thakurta, A. Smith, and A. Terzis. The Last Iterate Advantage: Empirical Auditing and Principled Heuristic Analysis of Differentially Private SGD. In *ICLR*, 2025.
- [85] NHS England. Purposes of the GP electronic health record. <https://www.england.nhs.uk/long-read/purpose-of-the-gp-electronic-health-record/>, 2025.
- [86] B. Niu, Z. Zhou, Y. Chen, J. Cao, and F. Li. DP-Opt: Identify High Differential Privacy Violation by Optimization. In *ICWASA*, 2022.
- [87] OpenDP. Developing Open Source Tools for Differential Privacy. <https://opendp.org/>, 2025.
- [88] A. Panda, X. Tang, M. Nasr, C. A. Choquette-Choo, and P. Mittal. Privacy Auditing of Large Language Models. In *ICLR*, 2025.
- [89] N. Papernot, M. Abadi, Ú. Erlingsson, I. Goodfellow, and K. Talwar. Semi-supervised Knowledge Transfer for Deep Learning from Private Training Data. In *ICLR*, 2017.
- [90] N. Papernot, P. McDaniel, A. Sinha, and M. P. Wellman. SoK: Security and Privacy in Machine Learning. In *IEEE S&P*, 2018.
- [91] K. Pillutla, G. Andrew, P. Kairouz, H. B. McMahan, A. Oprea, and S. Oh. Unleashing the Power of Randomization in Auditing Differentially Private ML. In *NeurIPS*, 2024.
- [92] A. Pyrgelis, C. Troncoso, and E. D. Cristofaro. Knock Knock, Who's There? Membership Inference on Aggregate Location Data. In *NDSS*, 2018.
- [93] J. Reed and B. C. Pierce. Distance makes the types grow stronger: a calculus for differential privacy. In *ICFP*, 2010.
- [94] A. Rényi. On measures of entropy and information. In *Berkeley Symposium on Mathematical Statistics and Probability*, 1961.
- [95] A. Roth. Lecture notes for “The Algorithmic Foundations of Data Privacy”. <https://www.cis.upenn.edu/~aaroth/courses/slides/Lecture11.pdf>, 2011.
- [96] A. Salem, G. Cherubin, D. Evans, B. Köpf, A. Paverd, A. Suri, S. Tople, and S. Zanella-Béguelin. SoK: Let the Privacy Games Begin! A Unified Treatment of Data Inference Privacy in Machine Learning. In *IEEE S&P*, 2023.
- [97] I. Sason and S. Verdú.  $f$ -divergence Inequalities. *IEEE Transactions on Information Theory*, 2016.
- [98] J. Schuchardt, M. Stoian, A. Kosmala, and S. Günemann. Unified Mechanism-Specific Amplification by Subsampling and Group Privacy Amplification. In *NeurIPS*, 2024.
- [99] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership Inference Attacks against Machine Learning Models. In *IEEE S&P*, 2017.
- [100] T. Steinke, M. Nasr, and M. Jagielski. Privacy Auditing with One (1) Training Run. In *NeurIPS*, 2024.
- [101] B. Stoddard, Y. Chen, and A. Machanavajjhala. Differentially Private Algorithms for Empirical Machine Learning. *arXiv:1411.5428*, 2014.
- [102] L. Sweeney. Simple demographics often identify people uniquely. *Health (San Francisco)*, 2000.

- [103] X. Tang, A. Panda, M. Nasr, S. Mahloujifar, and P. Mittal. Private Fine-tuning of Large Language Models with Zeroth-order Optimization. *TMLR*, 2025.
- [104] F. Tramer, A. Terzis, T. Steinke, S. Song, M. Jagielski, and N. Carlini. Debugging Differential Privacy: A Case Study for Privacy Auditing. *arXiv:2202.12219*, 2022.
- [105] Transport for London. Our open data. <https://tfl.gov.uk/info-for/open-data-users/our-open-data>, 2025.
- [106] US Census Bureau. United States Census. <https://www.census.gov/>, 2025.
- [107] H. Wang, Z. Xu, S. Jia, Y. Xia, and X. Zhang. Why current differential privacy schemes are inapplicable for correlated data publishing? In *WWW*, 2021.
- [108] Y. Wang, Z. Ding, D. Kifer, and D. Zhang. CheckDP: An Automated and Integrated Approach for Proving Differential Privacy or Finding Precise Counterexamples. In *CCS*, 2020.
- [109] Y. Wang, Z. Ding, G. Wang, D. Kifer, and D. Zhang. Proving differential privacy with shadow execution. In *PLDI*, 2019.
- [110] Z. Xiang, T. Wang, and D. Wang. Preserving Node-level Privacy in Graph Neural Networks. In *IEEE S&P*, 2024.
- [111] Z. Xiang, T. Wang, and D. Wang. Privacy Audit as Bits Transmission:(Im) possibilities for Audit by One Run. In *USENIX Security*, 2025.
- [112] J. Ye and R. Shokri. Differentially Private Learning Needs Hidden State (Or Much Faster Convergence). In *NeurIPS*, 2022.
- [113] S. Yeom, I. Giacomelli, A. Menaged, M. Fredrikson, and S. Jha. Overfitting, robustness, and malicious algorithms: A study of potential causes of privacy risk in machine learning. *Journal of Computer Security*, 2020.
- [114] S. Yoon, W. Jeung, and A. No. Optimizing Adversarial Samples for Tighter Privacy Auditing in Final Model-Only Settings. In *Statistical Frontiers in LLMs and Foundation Models Workshop NeurIPS*, 2024.
- [115] S. Zanella-Béguelin, L. Wutschitz, S. Tople, A. Salem, V. Rühle, A. Paverd, M. Naseri, B. Köpf, and D. Jones. Bayesian estimation of differential privacy. In *ICML*, 2023.
- [116] D. Zhang and D. Kifer. LightDP: Towards automating differential privacy proofs. In *POPL*, 2017.
- [117] J. Zhang, X. Xiao, and X. Xie. PrivTree: A Differentially Private Algorithm for Hierarchical Decompositions. In *SIGMOD*, 2016.
- [118] Y. Zhang, X. Liang, R. Du, and J. Tian. DP-Discriminator: A Differential Privacy Evaluation Tool Based on GAN. In *ICCR*, 2024.
- [119] Y. Zhu, J. Dong, and Y.-X. Wang. Optimal Accounting of Differential Privacy via Characteristic Function. In *AISTATS*, 2022.

## A Scopus Search Query

Our search query was set up to catch the following keywords in the title or abstract of the papers: *differential privacy, dif-*

```
(
  TITLE-ABS-KEY("differential privacy") OR
  TITLE-ABS-KEY("differentially private")
) AND
(
  TITLE-ABS-KEY("audit*") OR
  TITLE-ABS-KEY("find* violat*") OR
  TITLE-ABS-KEY("detect* violat*") OR
  TITLE-ABS-KEY("identify* violat*") OR
  TITLE-ABS-KEY("discover* violat*") OR
  TITLE-ABS-KEY("find* vuln*") OR
  TITLE-ABS-KEY("detect* vuln*") OR
  TITLE-ABS-KEY("identify* vuln*") OR
  TITLE-ABS-KEY("discover* vuln*") OR
  TITLE-ABS-KEY("find* privacy violat*") OR
  TITLE-ABS-KEY("detect* privacy violat*") OR
  TITLE-ABS-KEY("identify* privacy violat*") OR
  TITLE-ABS-KEY("discover* privacy violat*") OR
  TITLE-ABS-KEY("find* privacy vuln*") OR
  TITLE-ABS-KEY("detect* privacy vuln*") OR
  TITLE-ABS-KEY("identify* privacy vuln*") OR
  TITLE-ABS-KEY("discover* privacy vuln*") OR
  TITLE-ABS-KEY("violat* DP") OR
  TITLE-ABS-KEY("break* DP") OR
  TITLE-ABS-KEY("privacy estimat*") OR
  TITLE-ABS-KEY("estimat* privacy")
)
```

Figure 4: Scopus Query.

*ferentially private, audit, find/detect/discover privacy violations/vulnerabilities, break/violate DP, estimate privacy.* The exact search query we ran on the Scopus research database in January 2025 to collate the list of relevant prior work is provided in Figure 4.

## B Prior Work on DP auditing

In Table 5, we provide the full list of the 45 papers that form the basis of our systematization. Whereas, as discussed in Section 4, Table 1 only reports a selection of 21 papers (specifically, papers that are the first to audit a particular DP guarantee or mechanism, or that achieve the current “best” auditing results).

Reference	Mechanisms Audited	Foundations		Operational Details		Progress		
		DP Guarantee	Sub-Sampling	Attack	Eval. Function	Efficient	End-to-End	Tight
Ding et al. [34]	Report Noisy Max Noisy Hist SVT	$\epsilon$ -DP	-	-	Output-set	✗	✓	✓
Bischel et al. [18]	Report Noisy Max NoisySum AboveThreshold	$\epsilon$ -DP	-	-	Output-set	✗	✓	✓
Liu et al. [65]	Report Noisy Max Noisy Hist SVT Truncated Geometric	$(\epsilon, \delta)$ -DP	-	-	Output-set	✗	✓	✓
Jagielski et al. [54]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗	✓	✗
Bischel et al. [19]	Report Noisy Max Noisy Hist SVT RAPPOR Prefix Sum Truncated Geometric	$(\epsilon, \delta)$ -DP	-	DPD	Output-set	✗	✓	✓
Bernau et al. [17]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	DPD MI	Output-set	✗	✗	✓
Wang et al. [107]	Laplace	$\epsilon$ -DP	-	DPD	Output-set	✗	✓	✓
Nasr et al. [83]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗ ✗	✓ ✗	✗ ✓
Malek et al. [72]	PATE-FM ALIBI	Label DP	-	AI	Accuracy	✓	✓	✗
Lu et al. [67]	Naive Bayes Random Forest DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗	✓	✗
Houssiau et al. [52]	CTGAN MST PrivBayes	$(\epsilon, \delta)$ -DP	-	AI	FPR/FNR	✗	✓	✗
Niu et al. [86]	Report Noisy Max NoisyHist SVT Laplace RAPPOR	$(\epsilon, \delta)$ -DP	-	-	Output-set	✗	✓	✓
Tramèr et al. [104]	Backpropagation Clipping	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗	✓	✗
Askin et al. [12]	Report Noisy Max SVT Laplace Exponential Mechanism	$\epsilon$ -DP	-	-	Distance Est.	✗	✓	✓
Lokna et al. [66]	MST Gaussian Discrete Gaussian Laplace	$(\epsilon, \delta)$ -DP	-	DPD	Output-set	✗	✓	✓
Zanella-Béguelin et al. [115]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗	✓	✗
Nasr et al. [82]	DP-SGD	$f$ -DP	Poisson -	MI	FPR/FNR	✗ ✓	✓ ✗	✗ ✓
Gorla et al. [50]	Truncated Laplace	$(\alpha, \gamma)$ -RDP	-	-	Histogram Est.	✗	✓	✓
Maddock et al. [70]	DP-FedSGD	User-level DP	Poisson	MI	FPR/FNR	✓	✗	✗
Pillutla et al. [91]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	Custom	✗ ✗	✓ ✗	✗ ✗
Steinke et al. [100]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	Accuracy	✓ ✓	✓ ✗	✗ ✗
Galen et al. [4]	DP-FedAvg	User-level DP	Shuffle	MI	Variance Est.	✓	✗	✗
Matsumoto et al. [73]	LDP-SGD	$\epsilon$ -DP	Shuffle	MI	FPR/FNR	✗	✗	✗

Chadha et al. [21]	PATE CaPC PromptPATE Private k-NN	$(\alpha, \gamma)$ -RDP	–	MI	Custom	✗	✗	✗
Zhang et al. [118]	Report Noisy Max NoisyHist SVT Laplace RAPPOR Prefix Sum Truncated Geometric	$\epsilon$ -DP	–	DPD	Output-set	✗	✓	✓
Chida et al. [24]	PrivBayes DPCopula Randomized Response	MI	–	MI	FPR/FNR	✗	✓	✗
Lu et al. [68]	Report Noisy Max NoisyHist SVT	$(\alpha, \gamma)$ -RDP	–	DPD	Output-set	✗	✓	✓
Xiang et al. [110]	DP-GNN	$f$ -DP	Poisson	MI	FPRFNR	✓	✗	✓
Annamalai et al. [8]	PrivBayes CTGAN DP-WGAN	$(\epsilon, \delta)$ -DP	Poisson	MI	FPR/FNR	✗	✓	✗
Feng et al. [45]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	–	Output-set	✗	✓	✓
Annamalai et al. [5]	DP-SGD	$f$ -DP	Poisson	MI	FPRFNR	✗	✗	✓
Annamalai et al. [7]	DP-SGD	$\mu$ -GDP	–	MI	FPR/FNR	✗	✓	✗
Yoon et al. [114]	DP-SGD	$\mu$ -GDP	–	MI	FPR/FNR	✗	✓	✗
Kong et al. [62]	Laplace Gaussian SVT DP-SGD	$(\epsilon, \delta)$ -DP	–	–	Divergence Est.	✗	✗	✓
Debenedetti et al. [31]	DP-SGD	$f$ -DP	Poisson	MI	FPR/FNR	✗	✗	✗
Arcolezi et al. [10]	Generalized Randomized Response Subset Selection Local Hashing Unary Encoding Histogram Encoding	$(\epsilon, \delta)$ -DP	–	DPD	Output-set	✗	✓	✓
Mahloujifar et al. [71]	DP-SGD Gaussian	$f$ -DP $\mu$ -GDP	Poisson –	Recon	Accuracy	✓ ✓	✗ ✓	✗ ✓
Annamalai et al. [6]	DP-SGD (Shuffle)	$(\epsilon, \delta)$ -DP	Shuffle	MI	FPR/FNR	✗	✗	✗
Huang et al. [53]	Laplace Gaussian DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	–	Variance Est.	✗	✗	✓
Nasr et al. [84]	DP-SGD	$f$ -DP	Poisson	MI	FPR/FNR	✗	✓	✗
Cebere et al. [20]	DP-SGD	$\mu$ -GDP	Poisson	MI	FPR/FNR	✗	✓	✗
Ganev et al. [47]	PATE-GAN	$(\epsilon, \delta)$ -DP	–	MI	FPR/FNR	✗	✓	✗
Askin et al. [11]	Laplace Gaussian Sub-Sampling DP-SGD	$f$ -DP	Poisson	DPD	Output-set	✗	✓	✗
Panda et al. [88]	DP-SGD	$(\epsilon, \delta)$ -DP	Poisson	MI	Accuracy	✓	✓	✗
Xiang et al. [111]	DP-SGD Gaussian	$f$ -DP $\mu$ -GDP	Poisson –	MI	Accuracy	✓ ✓	✗ ✓	✗ ✓

**Table 5:** Full summary of key prior work on DP auditing.